



Machine Learning for Slow but Steady Interplanetary Construction


Adrian Agogino



August 29th, 2017

Problem Statement

- Can a set of construction robots build useful structures on other planets given a long period of time
- Large number of robots (dozens)
- Robots are small
- Construction time up to 20 years
- Simple structures:
 - Landing pads
 - Pyramid type buildings



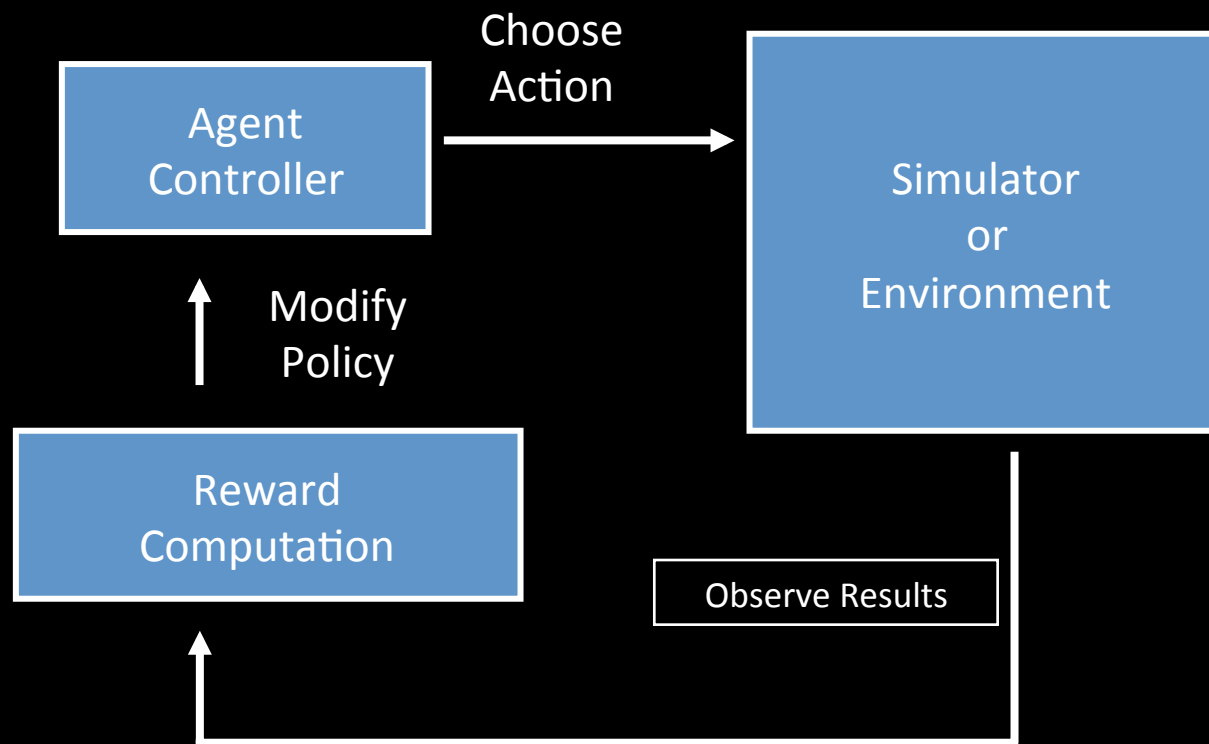
Challenges

- Robots need to coordinate
- Robots need to cooperate
 - Not large enough to operate entirely in parallel
- Robot precision is low
- Raw material is poor
- Robots need to overcome failure
- Robots need to be autonomous


Multiagent Learning System

- Allows for coordination of multiple agents
- Assumes agents are imperfect and low precision
- Adapts to agent failures
- Handles uncertain environments
- Learns instead of being told what to do
- (Can be slow)

Agent Learning





Temporal Credit Assignment



Temporal Credit Assignment Problem
(RL)


Structural Credit Assignment




Both Credit Assignment Problems

- Can be difficult to assign credit through time and to the right agent.
- Convergence possible through “team games” but very slow.
- Reward shaping can improve prospects
 - Want shaped reward to improve “signal.”
 - Want shaped reward to promote coordination.

Control Problem



Evaluation Problem



Network 1
Network 2
Network 3


Rover updates its population of mapping functions based on evaluation

Shaped evaluation is used to take "noise" out of global evaluation

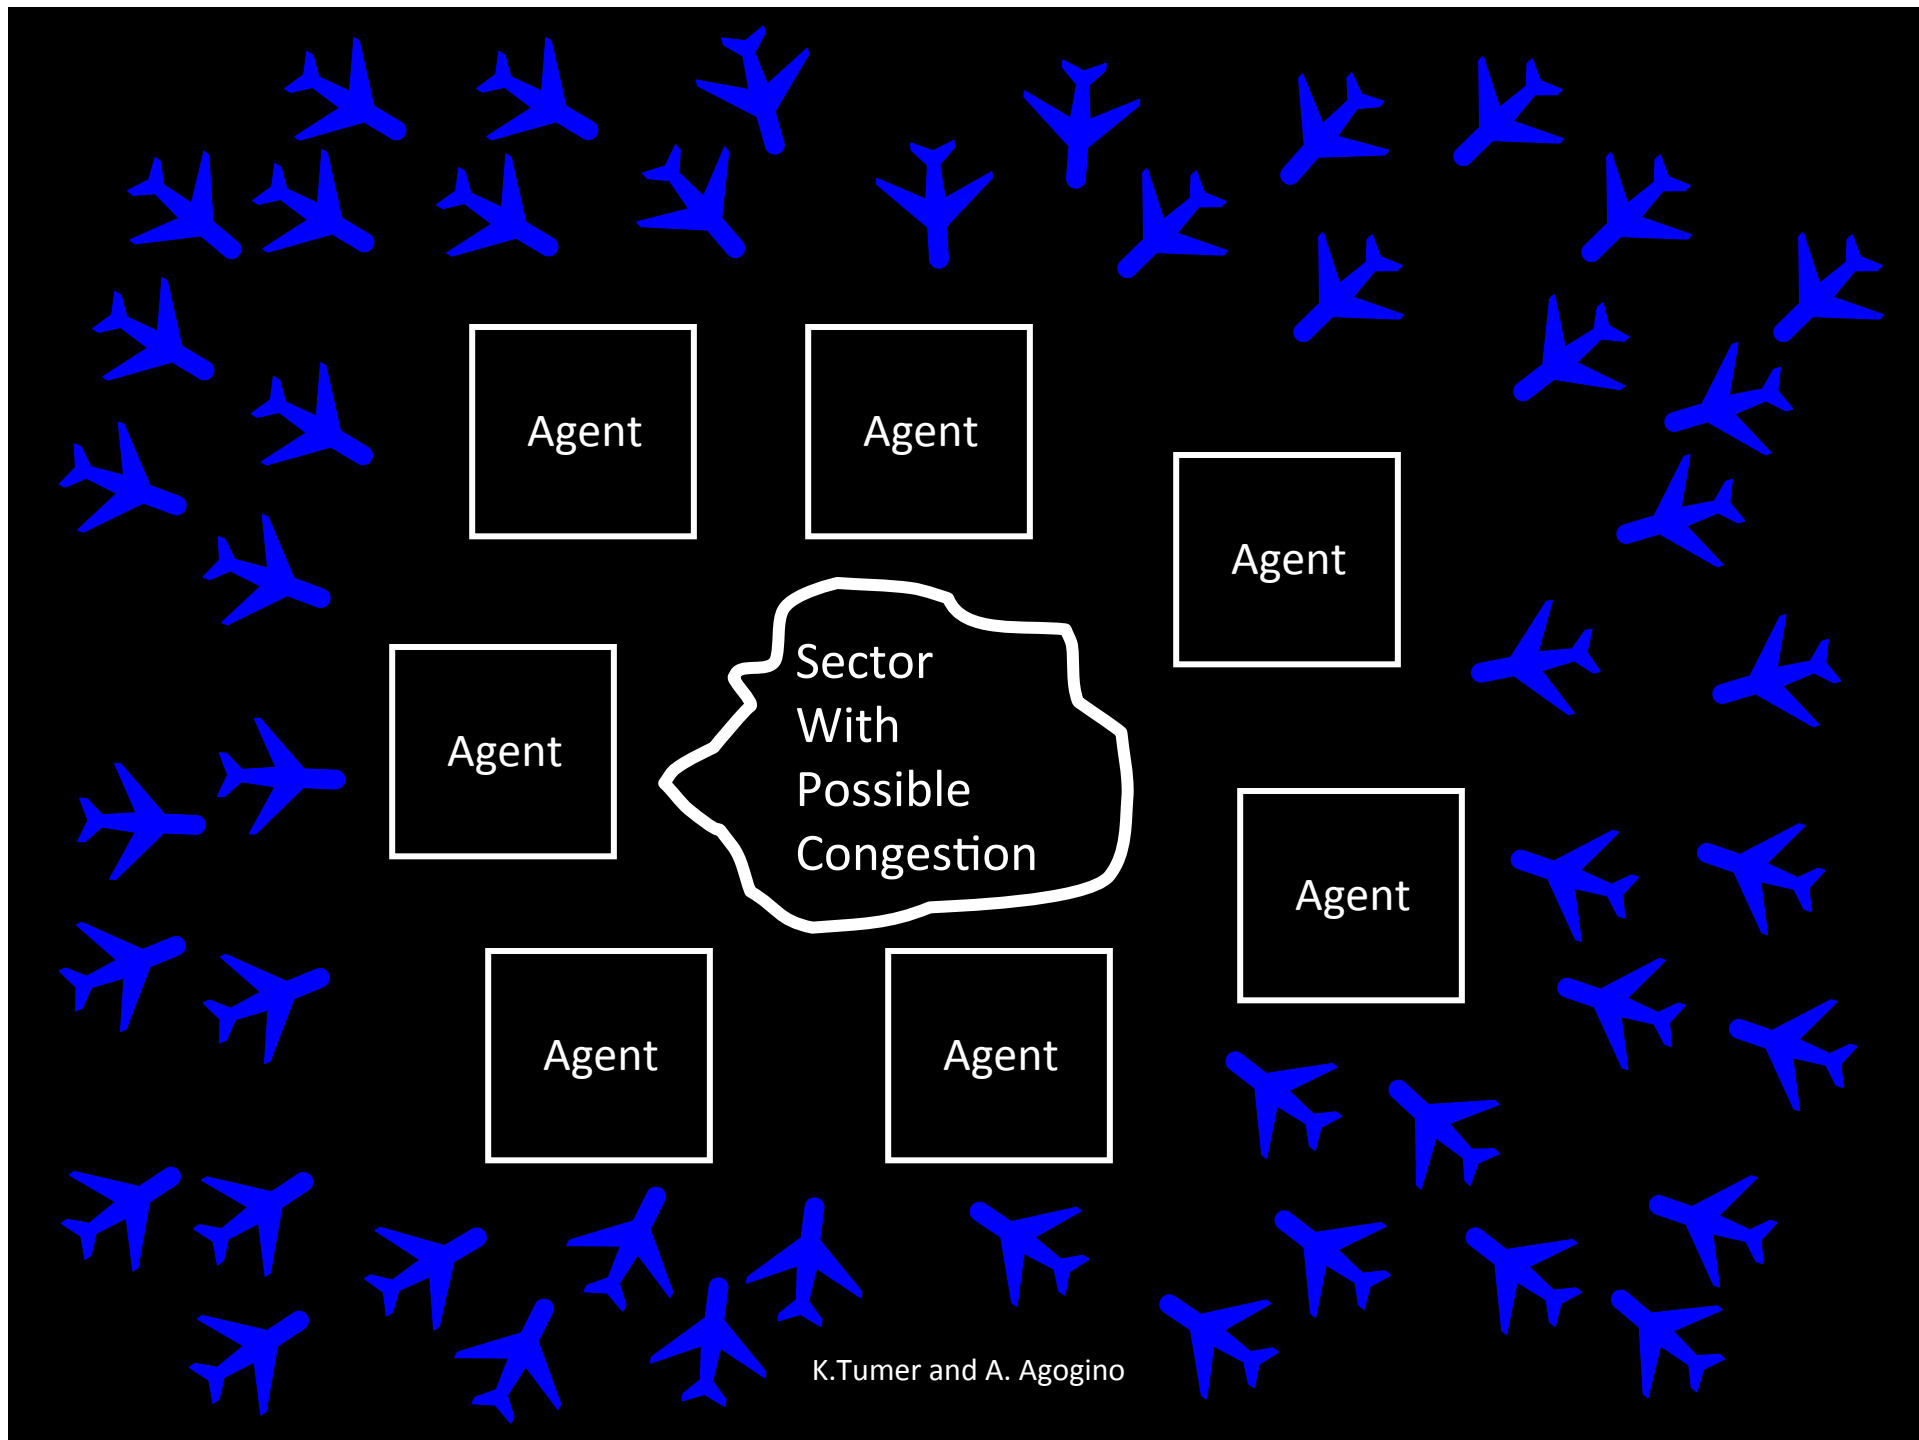
$$D_{\eta} \equiv G(z) - G(z_{-\eta} + c_{\eta})$$

State of system

State of system without rover




Rover is evaluated based on contribution to global evaluation






Will it Work for Planetary Construction?

- Is enough coordination possible?
- Can learning be independent of operators?
- Will small rovers have enough capabilities?



Tensegrity Robots







In Action



Tensegrity Control Learning





Tensegrity Construction



Will it Work?

- Collapsible and light weight
- Scalable
- Can be highly redundant
- (Learning can be difficult)



Thank You

